

CLAIMS

What is claimed is:

1. A method for balancing transmission unit traffic over
5 network links, comprising:

a. disposing transmission units into flows;
b. grouping flows into first flow lists, each of the first
flow lists corresponding to a selected network link;
c. determining a traffic metric representative of a traffic
10 load on the selected network link;

d. responsive to the traffic metric, regrouping flows into
second flow lists corresponding to the selected network link, the
regrouping balancing the transmission unit traffic among the
network links; and

15 e. transmitting the respective second flow list over the
respective selected network link.

2. The method of Claim 1, wherein each of the transmission
units comprise one of a packet, a frame, a cell, and a combination
20 thereof.

3. The method of Claim 1, wherein each of the transmission
units includes one of source information, destination information,
and a combination thereof, and the disposing comprises
25 characterizing each of the transmission units according to one of
the source information, the destination information, and a
combination thereof.

4. The method of Claim 1, wherein the transmission unit traffic is effected using a predetermined link-layer transmission protocol.

5. The method of Claim 4, wherein the predetermined link-layer transmission protocol is one of a packet-based protocol, a cell-based protocol, a frame-based protocol, and a combination thereof.

6. The method of Claim 5, wherein the predetermined link-layer transmission protocol is one of an IEEE STD. 802 protocol, an ATM protocol, an FDDI protocol, an X.25 protocol, an ISDN protocol, a CDPD protocol, and a Frame Relay protocol.

7. The method of Claim 4, wherein the predetermined link-layer transmission protocol communicates the transmission unit traffic over the network links in cooperation with a network-layer protocol.

8. The method of Claim 4, wherein the predetermined link-layer transmission protocol communicates the transmission unit traffic over the network links in cooperation with a transport-layer protocol.

9. The method of Claim 7, wherein the network-layer protocol cooperates with a transport-layer protocol to communicate the transmission unit traffic across the network links.

10. The method of Claim 7, wherein the network-layer protocol is one of a connectionless protocol and a connection-based protocol.

5 11. The method of Claim 10, wherein the network-layer protocol is a connectionless protocol, and the connectionless protocol is an internet protocol (IP).

10 12. The method of Claim 10, wherein the network-layer protocol is one of an internet protocol (IP), a Switched Multi-megabit Data Service protocol (SMDS), a general packet radio service (GPRS), a Message Transfer Part Level 3 protocol (MTP3), a an internet packet exchange protocol (IPX), an X.25/X.75 protocol, a connectionless network protocol (CLNP), internet datagram protocol (IDP), a datagram delivery protocol (DDP), a Xerox Network Systems protocol (XNS), and a combination thereof.

15 13. The method of Claim 9, wherein the transport-layer protocol is one of a connectionless protocol and a connection-based protocol.

20 14. The method of Claim 13, wherein the transport-layer protocol is a connection-based protocol, and the connection-based protocol is a transmission control protocol (TCP).

25 15. The method of Claim 13, wherein the transport-layer protocol is one of a transmission control protocol (TCP), a user datagram protocol (UDP), a NetBIOS protocol, an H.323 protocol, a GSM/CDMA protocol, a Stream Control Transmission Protocol (SCTP),
30 and a combination thereof.

16. The method of Claim 9, wherein the network-layer protocol is a connectionless protocol and the connectionless protocol is an internet protocol (IP), and wherein the transport-layer protocol is a connection-based protocol, and the connection-based protocol is a transmission control protocol (TCP).

17. The method of Claim 1, wherein the network links are network links having heterogeneous speed.

18. The method of Claim 3, wherein the first flow lists are ordered flow lists.

19. The method of Claim 18, wherein each of the ordered flow list comprises a linked list ordered by flow size.

20. The method of Claim 19, wherein the ordered flow list is ordered by decreasing flow size.

21. The method of Claim 1, wherein determining a Traffic Metric representative of a traffic load on the selected network link comprises:

$$T_i(f_x, t) = \frac{\frac{N_{Li}(f_x, t) * K_i + N_{Si}(f_x, t)}{\Delta t} + T_i(f_x, t-1)}{2}$$

where

$T_i(f_x, t)$ is a Traffic Metric of Flow f_x in selected

link i , sampled at time t ;

$N_{Li}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size greater than or equal to a preselected size threshold, observed between time $t - 1$ and time t ;

5 $N_{Si}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size less than a preselected size threshold, observed between sampling time $t - 1$ and time t ;

K_{ix} is a predetermined link load factor; and

10 Δt is the inter-sampling time, measured as the interval between time $t - 1$ and time t , in selected link i .

22. The method of Claim 21, wherein the Traffic Metric
15 comprises an Aggregate Traffic Metric characterized by:

$$TA(i, t) = \sum_x T_i(f_x, t)$$

where

$TA(i, t)$ is an Aggregate Traffic Metric for selected link i , and

20 $T_i(f_x, t)$ is the Traffic Metric of Flow f_x in link i , sampled at time t .

23. The method of Claim 22, wherein the Traffic Metric comprises a Scaled Aggregate Traffic Metric characterized by:

$$STA(i,t) = TA(i,t) * S(i,t)$$

5 where

$STA(i,t)$ is the Scaled Aggregated Traffic Metrics observed at time t for network link i ;

$TA(i,t)$ is the Traffic Metric observed at time t for network link i ; and

10 $S(i,t)$ is a preselected scaling factor for network link i at time t .

24. The method of Claim 23, wherein the Traffic Metric comprises a Link Group Arithmetic Mean characterized by:

$$M(t) = \frac{\sum_{i=1}^n STA(i,t)}{n}$$

15 where

$M(t)$ is the Link Group Arithmetic Mean of Scaled Aggregated Traffic Metrics summed over all network links i ;

20 $STA(i,t)$ is the Scaled Aggregated Traffic Metrics observed at time t for network link i ; and

n is the number of active network links.

25 25. The method of Claim 24, wherein the Traffic Metric comprises a Link Group Absolute Deviation characterized by:

$$D(t) = \frac{\sum_{i=1}^n |STA(i,t) - M(t)|}{n}$$

where

$D(t)$ is the Link Group Absolute Deviation of the Scaled Aggregated Traffic Metrics summed over all network links i ;

$STA(i,t)$ is the Scaled Aggregated Traffic Metric observed at time t for network link i ;

$M(t)$ is the Link Group Arithmetic Mean of the Scaled Aggregated Traffic Metrics summed over all network links i ; and

n is the number of active network links.

26. The method of Claim 25, wherein:

a. a HIGH traffic load on link i is characterized by

$$STA(i,t) \geq M(t) + D(t);$$

b. a LOW traffic load on link i is characterized by

$$STA(i,t) < M(t) - D(t); \text{ and}$$

c. a NORMAL traffic load on link i is characterized by

$$M(t) + D(t) > STA(i,t) \geq M(t) - D(t).$$

27. The method of Claim 26, wherein the regrouping comprises moving a preselected flow from a first flow list corresponding to a first preselected link having a HIGH traffic load to a second flow list corresponding to a second preselected link having one of a NORMAL traffic load and a LOW traffic load.

28. The method of Claim 21, wherein the predetermined link load factor, K_{ix} , for Flow x in link i is between about 1.0 and about 10.0.

5 29. The method of Claim 28, wherein the predetermined link load factor, K_{ix} , for Flow x in link i is between about 4.0 and about 5.0.

10 30. A method for balancing transmission unit traffic over heterogeneous speed network links, comprising:

15 a. disposing transmission units into flows, wherein each of the transmission units includes one of source information, destination information, and a combination thereof, and the disposing comprises characterizing each of the transmission units according to one of the source information, the destination information, and a combination thereof, and wherein each of the transmission units comprise one of a packet, a frame, a cell, and a combination thereof;

20 b. grouping flows into first flow lists, each of the first decreasing-size-ordered linked flow lists corresponding to a selected network link;

 c. determining a traffic metric representative of a traffic load on the selected network link;

25 d. responsive to the traffic metric, regrouping flows into second decreasing-size-ordered linked flow lists corresponding to the selected network link, the regrouping balancing the transmission unit traffic among the network links; and

30 e. transmitting the respective second flow list over the respective selected network link using a predetermined link-layer transmission protocol, wherein the predetermined link-layer

transmission protocol communicates the transmission unit traffic over the network links in cooperation with a network-layer protocol; wherein the network-layer protocol cooperates with a transport-layer protocol to communicate the transmission unit traffic across the network links, and wherein each of the network-layer protocol and the transport-layer protocol is one of a connectionless protocol and a connection-based protocol.

31. The method of Claim 30, wherein the predetermined link-layer transmission protocol is one of an IEEE STD. 802 protocol, an ATM protocol, an FDDI protocol, an X.25 protocol, an ISDN protocol, a CDPD protocol, and a Frame Relay protocol; wherein the network-layer protocol is one of an internet protocol (IP), a Switched Multi-megabit Data Service protocol (SMDS), a general packet radio service (GPRS), a Message Transfer Part Level 3 protocol (MTP3), a an internet packet exchange protocol (IPX), an X.25/X.75 protocol, a connectionless network protocol (CLNP), internet datagram protocol (IDP), a datagram delivery protocol (DDP), a Xerox Network Systems protocol (XNS), and a combination thereof; and wherein the transport-layer protocol is one of a transmission control protocol (TCP), a user datagram protocol (UDP), a NetBIOS protocol, an H.323 protocol, a GSM/CDMA protocol, a Stream Control Transmission Protocol (SCTP), and a combination thereof.

32. The method of Claim 30, wherein the predetermined link-layer transmission protocol is an IEEE STD. 802.3 protocol, wherein the network-layer protocol is an internet protocol (IP), and wherein the transport-layer protocol is a transmission control protocol (TCP).

33. The method of Claim 32, wherein the Traffic Metric representative of a traffic load on the selected network link comprises a Link Group Arithmetic Mean and the determining the Traffic Metric comprises:

- 5 a. determining a Link Traffic Flow Metric characterized by

$$T_i(f_x, t) = \frac{\frac{N_{Li}(f_x, t) * K_i + N_{Si}(f_x, t)}{\Delta t} + T_i(f_x, t-1)}{2}$$

where

$T_i(f_x, t)$ is a Traffic Metric of Flow f_x in selected link i , sampled at time t ,

10 $N_{Li}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size greater than or equal to a preselected size threshold, observed between time $t - 1$ and time t ,

15 $N_{Si}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size less than a preselected size threshold, observed between sampling time $t - 1$ and time t ,

20 K_{ix} is a predetermined link load factor for Flow x in link i having a value in a range between about 4.0 and about 5.0, and

Δt is the inter-sampling time, measured as the interval between time $t - 1$ and time t , in selected link i ;

- b. determining an Aggregate Traffic Metric characterized by

$$TA(i,t) = \sum_x T_i(f_x,t)$$

where

$TA(i,t)$ is an Aggregate Traffic Metric for selected
link i , and

$T_i(f_x,t)$ is the Traffic Metric of Flow f_x in link i ,
sampled at time t ;

- c. determining a Scaled Aggregate Traffic Metric
characterized by

$$STA(i,t) = TA(i,t) * S(i,t)$$

where

$STA(i,t)$ is the Scaled Aggregated Traffic Metrics
observed at time t for network link i ,

$TA(i,t)$ is the Traffic Metric observed at time t for
network link i , and

$S(i,t)$ is a preselected scaling factor for network
link i at time t ; and

- d. determining a Link Group Arithmetic Mean characterized by

$$M(t) = \frac{\sum_{i=1}^n STA(i,t)}{n}$$

where

$M(t)$ is the Link Group Arithmetic Mean of Scaled
Aggregated Traffic Metrics summed over all network links
 i ,

$STA(i,t)$ is the Scaled Aggregated Traffic Metrics observed at time t for network link i , and n is the number of active network links.

34. The method of Claim 33, further comprising determining a Link Group Absolute Deviation characterized by

$$D(t) = \frac{\sum_{i=1}^n |STA(i,t) - M(t)|}{n}$$

where

$D(t)$ is the Link Group Absolute Deviation of the Scaled Aggregated Traffic Metrics summed over all network links i ,

$STA(i,t)$ is the Scaled Aggregated Traffic Metric observed at time t for network link i ,

$M(t)$ is the Link Group Arithmetic Mean of the Scaled Aggregated Traffic Metrics summed over all network links i , and

n is the number of active network links; and

b. wherein

(1) a HIGH traffic load on link i is characterized by $STA(i,t) \geq M(t) + D(t)$,

(2) a LOW traffic load on link i is characterized by $STA(i,t) < M(t) - D(t)$, and

(3) a NORMAL traffic load on link i is characterized by $M(t) + D(t) > STA(i,t) \geq M(t) - D(t)$.

35. The method of Claim 34, wherein the regrouping comprises moving a preselected flow from a first flow list corresponding to a first preselected link having a HIGH traffic load to a second flow list corresponding to a second preselected link having one of a NORMAL traffic load and a LOW traffic load.

36. A method for transmitting transmission units through a network, comprising:

a. receiving a transmission unit from a transmission unit source;

b. classifying the transmission unit according to a predetermined flow characteristic;

c. selecting a preselected network link over which the transmission unit is to be transmitted; and

d. transmitting the transmission unit over the preselected network link.

37. The method of Claim 36, further comprising selecting the preselected network link according to the predetermined flow characteristic using a predetermined dynamic load balancing technique.

38. The method of Claim 37, further comprising monitoring the operation of a plurality of preselected network links, and re-assigning the predetermined flow characteristic from a first preselected network link to a second preselected network link, if the first preselected network link operationally fails.

39. A computer program product recorded on a computer readable medium for balancing transmission unit traffic over network links, comprising:

a. computer readable program code which disposes
5 transmission units into flows;

b. computer readable program code which groups flows into first flow lists, each of the first flow lists corresponding to a selected network link;

c. computer readable program code which determines a traffic
10 metric representative of a traffic load on the selected network link;

d. computer readable program code which, responsive to the traffic metric, re-assigns flows into second flow lists corresponding to the selected network link, the re-assigning
15 balancing the transmission unit traffic among the network links; and

e. computer readable program code which transmits the respective second flow list over the respective selected network link.

40. The computer program product of Claim 39, wherein each of the transmission units comprises one of a packet, a frame, a cell, and a combination thereof.

41. The computer program product of Claim 39, wherein each of the transmission units includes one of source information, destination information, and a combination thereof, and the disposing comprises characterizing each of the transmission units according to one of the source information, the destination
30 information, and a combination thereof.

42. The computer program product of Claim 39, wherein the transmission unit traffic is effected using a predetermined link-layer transmission protocol.

5 43. The computer program product of Claim 42, wherein the predetermined link-layer transmission protocol is one of a packet-based protocol, a cell-based protocol, a frame-based protocol, and a combination thereof.

10 44. The computer program product of Claim 43, wherein the predetermined link-layer transmission protocol is one of an IEEE STD. 802 protocol, an ATM protocol, an FDDI protocol, an X.25 protocol, an ISDN protocol, a CDPD protocol, and a Frame Relay protocol.

15 45. The computer program product of Claim 42, wherein the predetermined link-layer transmission protocol communicates the transmission unit traffic over the network links in cooperation with a network-layer protocol.

20 46. The computer program product of Claim 42, wherein the predetermined link-layer transmission protocol communicates the transmission unit traffic over the network links in cooperation with a transport-layer protocol.

25 47. The computer program product of Claim 45, wherein the network-layer protocol cooperates with a transport-layer protocol to communicate the transmission unit traffic across the network links.

48. The computer program product of Claim 45, wherein the network-layer protocol is one of a connectionless protocol and a connection-based protocol.

5 49. The computer program product of Claim 48, wherein the network-layer protocol is a connectionless protocol, and the connectionless protocol is an internet protocol (IP).

10 50. The computer program product of Claim 48, wherein the network-layer protocol is one of an internet protocol (IP), a Switched Multi-megabit Data Service protocol (SMDS), a general packet radio service (GPRS), a Message Transfer Part Level 3 protocol (MTP3), a an internet packet exchange protocol (IPX), an X.25/X.75 protocol, a connectionless network protocol (CLNP),
15 internet datagram protocol (IDP), a datagram delivery protocol (DDP), a Xerox Network Systems protocol (XNS), and a combination thereof.

20 51. The computer program product of Claim 47, wherein the transport-layer protocol is one of a connectionless protocol and a connection-based protocol.

25 52. The computer program product of Claim 51, wherein the transport-layer protocol is a connection-based protocol, and the connection-based protocol is a transmission control protocol (TCP).

53. The computer program product of Claim 51, wherein the transport-layer protocol is one of a transmission control protocol (TCP), a user datagram protocol (UDP), a NetBIOS protocol, an H.323 protocol, a GSM/CDMA protocol, a Stream Control Transmission Protocol (SCTP), and a combination thereof.

54. The computer program product of Claim 47, wherein the network-layer protocol is a connectionless protocol and the connectionless protocol is an internet protocol (IP), and wherein the transport-layer protocol is a connection-based protocol, and the connection-based protocol is a transmission control protocol (TCP).

55. The computer program product of Claim 39, wherein the network links are network links having heterogeneous speed.

56. The computer program product of Claim 41, wherein the first flow lists are ordered flow lists.

57. The computer program product of Claim 56, wherein each of the ordered flow list comprises a linked list ordered by flow size.

58. The computer program product of Claim 57, wherein the ordered flow list is ordered by decreasing flow size.

59. The computer program product of Claim 39, wherein determining a Traffic Metric representative of a traffic load on the selected network link comprises:

$$T_i(f_x, t) = \frac{N_{L_i}(f_x, t) * K_i + N_{S_i}(f_x, t)}{\Delta t} + T_i(f_x, t-1)$$

5 where

$T_i(f_x, t)$ is a Traffic Metric of Flow f_x in selected link i , sampled at time t ;

$N_{L_i}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size greater than or equal to a preselected size threshold, observed between time $t - 1$ and time t ;

$N_{S_i}(f_x, t)$ is the number of transmission units of Flow f_x in selected link i , with a transmission unit size less than a preselected size threshold, observed between sampling time $t - 1$ and time t ;

K_{ix} is a predetermined link load factor; and

Δt is the inter-sampling time, measured as the interval between time $t - 1$ and time t , in selected link i .

60. The computer program product of Claim 59, wherein the Traffic Metric comprises an Aggregate Traffic Metric characterized by:

$$TA(i, t) = \sum_x T_i(f_x, t)$$

where

$TA(i, t)$ is an Aggregate Traffic Metric for selected link i , and

$T_l(f_x, t)$ is the Traffic Metric of Flow f_x in link i ,
5 sampled at time t .

61. The computer program product of Claim 60, wherein the Traffic Metric comprises a Scaled Aggregate Traffic Metric characterized by:

$$STA(i, t) = TA(i, t) * S(i, t)$$

where

$STA(i, t)$ is the Scaled Aggregated Traffic Metrics observed at time t for network link i ;

$TA(i, t)$ is the Traffic Metric observed at time t for network link i ; and

$S(i, t)$ is a preselected scaling factor for network link i at time t .

62. The computer program product of Claim 61, wherein the Traffic Metric comprises a Link Group Arithmetic Mean characterized by:

$$M(t) = \frac{\sum_{i=1}^n STA(i, t)}{n}$$

where

$M(t)$ is the Link Group Arithmetic Mean of Scaled Aggregated Traffic Metrics summed over all network links i ;

$STA(i,t)$ is the Scaled Aggregated Traffic Metrics observed at time t for network link i ; and

n is the number of active network links.

63. The computer program product of Claim 62, wherein the Traffic Metric comprises a Link Group Absolute Deviation characterized by:

$$D(t) = \frac{\sum_{i=1}^n |STA(i,t) - M(t)|}{n}$$

where

$D(t)$ is the Link Group Absolute Deviation of the Scaled Aggregated Traffic Metrics summed over all network links i ;

$STA(i,t)$ is the Scaled Aggregated Traffic Metric observed at time t for network link i ;

$M(t)$ is the Link Group Arithmetic Mean of the Scaled Aggregated Traffic Metrics summed over all network links i ; and

n is the number of active network links.

64. The computer program product of Claim 63, wherein:
a. a HIGH traffic load on link i is characterized by

$$STA(i,t) \geq M(t) + D(t) ;$$

b. a LOW traffic load on link *i* is characterized by

$$STA(i,t) < M(t) - D(t) ; \text{ and}$$

c. a NORMAL traffic load on link *i* is characterized by

$$M(t) + D(t) > STA(i,t) \geq M(t) - D(t) .$$

65. The computer program product of Claim 64, wherein the regrouping comprises moving a preselected flow from a first flow list corresponding to a first preselected link having a HIGH traffic load to a second flow list corresponding to a second preselected link having one of a NORMAL traffic load and a LOW traffic load.

66. The computer program product of Claim 59, wherein the predetermined link load factor, K_{ix} , for Flow *x* in link *i* is between about 1.0 and about 10.0.

67. The computer program product of Claim 66, wherein the predetermined link load factor, K_{ix} , for Flow *x* in link *i* is between about 4.0 and about 5.0.

68. A network load balancer in a communication network having network links, comprising:

a. a flow synthesizer that receives transmission units from a transmission unit source, and synthesizes flows characteristic of selected transmission units; and

b. a link classifier, coupled with the flow synthesizer and the network links, that classifies the network links relative to a

predetermined flow metric, and assigns selected flows to selected network links responsive to the predetermined flow metric, the selected transmission units corresponding to the selected flows being communicated with the communication network through the respective selected network links.

69. The network load balancer of Claim 68, further comprising a flow classifier, coupled with the link classifier, that determines flow characteristics representative of a plurality of selected flows, and responsive thereto, assigns the plurality of selected flows to a selected link such that the selected transmission units respectively represented by the plurality of selected flows are communicated with the communication network through the respective selected network links.

70. The network load balancer of Claim 69, further comprising a link failure detector that causes the link classifier to re-assign the selected flows to others of the selected links in response to a network link failure.